

Data series subtraction with unknown and unmodeled background noise

Stefano Vitale,^{*} Giuseppe Congedo,[†] Rita Dolei, Valerio Ferroni, Mauro Hueller,
Daniele Vetrugno, and William Joseph Weber

*Department of Physics, University of Trento, and Trento Institute for Fundamental Physics
and Applications, INFN, 38123 Povo, Trento, Italy*

Heather Audley, Karsten Danzmann, Ingo Diepholz, Martin Hewitson, and Natalia Korsakova
*Albert-Einstein-Institut, Max-Planck-Institut für Gravitationsphysik und Universität Hannover,
Callinstrasse 38, 30167 Hannover, Germany*

Luigi Ferraioli
ETH Zürich, Institut für Geophysik, Sonneggstrasse 5, 8092 Zürich, Switzerland

Ferran Gibert, Nikolaos Karnesis, and Miquel Nofrarias
*Institut de Ciències de l'Espai, (CSIC-IEEC), Facultat de Ciències, Campus UAB, Torre C-5,
08193 Bellaterra, Spain*

Henri Inchauspe and Eric Plagnol
*APC, Université Paris Diderot, CNRS/IN2P3, CEA/Ifnu, Observatoire de Paris, Sorbonne Paris Cité,
10 Rue A. Domon et L. Duquet, 75205 Paris Cedex 13, France*

Oliver Jennrich and Paul W. McNamara
*European Space Technology Centre, European Space Agency, Keplerlaan 1, 2200 AG Noordwijk,
The Netherlands*

Michele Armano
*European Space Astronomy Centre, European Space Agency, Camino bajo del Castillo s/n,
Urbanización Villafranca del Castillo, Villanueva de la Cañada, 28692 Madrid, Spain*

James Ira Thorpe
*Gravitational Astrophysics Laboratory, NASA Goddard Space Flight Center, Greenbelt,
Maryland 20771, USA*

Peter Wass
*High Energy Physics Group, Imperial College London, Blackett Laboratory, Prince Consort Road,
London SW7 2AZ, United Kingdom*

(Received 18 April 2014; published 11 August 2014)

LISA Pathfinder (LPF), the precursor mission to a gravitational wave observatory of the European Space Agency, will measure the degree to which two test masses can be put into free fall, aiming to demonstrate a suppression of disturbance forces corresponding to a residual relative acceleration with a power spectral density (PSD) below $(30 \text{ fm/s}^2/\sqrt{\text{Hz}})^2$ around 1 mHz. In LPF data analysis, the disturbance forces are obtained as the difference between the acceleration data and a linear combination of other measured data series. In many circumstances, the coefficients for this linear combination are obtained by fitting these data series to the acceleration, and the disturbance forces appear then as the data series of the residuals of the fit. Thus the background noise or, more precisely, its PSD, whose knowledge is needed to build up the likelihood function in ordinary maximum likelihood fitting, is here unknown, and its estimate constitutes instead one of the goals of the fit. In this paper we present a fitting method that does not require the knowledge of the PSD of the background noise. The method is based on the analytical marginalization of the posterior parameter probability density with respect to the background noise PSD, and returns an estimate both for the fitting parameters and for the PSD. We show that both these estimates are unbiased, and that, when using averaged Welch's periodograms for the residuals, the estimate of the PSD is consistent, as its error tends to zero with the inverse square root of the number of averaged periodograms. Additionally, we find that the method is equivalent to some

^{*}stefano.vitale@unitn.it

[†]Present address: Department of Physics, University of Oxford, Oxford OX1 3RH, UK.

implementations of iteratively reweighted least-squares fitting. We have tested the method both on simulated data of known PSD and on data from several experiments performed with the LISA Pathfinder end-to-end mission simulator.

DOI: [10.1103/PhysRevD.90.042003](https://doi.org/10.1103/PhysRevD.90.042003)

PACS numbers: 04.80.Nn, 07.05.Kf, 95.55.-n

I. INTRODUCTION

LISA Pathfinder (LPF) [1] is the precursor mission to a gravitational wave (GW) observatory of the European Space Agency (ESA). Its primary goal is that of assessing if a set of reference test masses (TMs) can be put into free motion, with residual accelerations, relative to the local inertial frame, having a power spectral density (PSD) less than $(30 \text{ fm/s}^2/\sqrt{\text{Hz}})^2$, at frequencies between 1 and 30 mHz. This goal is pursued by measuring the relative acceleration of two TMs, separated by a nominal distance of 38 cm, along the line—whose direction we call x —joining their centres of mass (Fig. 1). The relative motion between the TMs, x_{12} , is measured by means of a laser interferometer, the output of which $s_{12} = x_{12} + n_{12}$ is affected by a readout noise n_{12} with less than $(6 \text{ pm}/\sqrt{\text{Hz}})^2$ PSD at mHz frequencies.

The relative acceleration a is then calculated by numerically performing the second time derivative [2] of the interferometer output s_{12} ,




FIG. 1 (color online). Schematic of LPF. The figure shows the reference TM, TM2, and the two laser interferometers—represented by their respective laser beam paths—that measure x_1 and x_{12} , respectively. The x axis, shown in the figure, is parallel to the line joining the centres of mass of the two TMs. The z axis, normal to the figure, points toward the Sun. Also shown are the electrodes used to apply the forces to TM2, necessary to keep it at nominally fixed distance from the reference TM. Similarly, the picture shows a pair of μN thrusters that are used to force the satellite to stay at a nominally fixed position relative to the reference TM. Not shown in the figure are the electrodes and the μN thrusters used to control TM and satellite along degrees of freedom other than x .

$$a = \frac{d^2 s_{12}}{dt^2} = \frac{d^2 x_{12}}{dt^2} + \frac{d^2 n_{12}}{dt^2} \equiv \frac{d^2 x_{12}}{dt^2} + a_r, \quad (1)$$

where we have implicitly defined the readout acceleration noise a_r .

The TMs are not both free-falling along x . One TM, the inertial reference, is indeed following a pure geodesic orbit, but both the satellite, and the other TM, that we call TM2, are forced, by some force control loop, to stay nominally at fixed positions relative to the reference TM. The satellite is actuated by a set of μN thrusters within a feedback loop driven by the signal from a dedicated interferometer, which measures the relative displacement x_1 between the satellite and the reference TM. The second TM is instead subject to a weak electrostatic force commanded by a feedback loop driven by the main interferometer signal s_{12} .

The relative motion of the satellite and the TMs, along degrees of freedom other than x , is also measured, either by laser interferometers or by capacitive sensing, and controlled by a combination of electrostatic forces and torques on the TMs, and of μN -thruster-generated forces and torques acting on the satellite.

In standard operations, control loops keep the relative motion small enough that the system is expected to behave linearly, obeying a set of linear dynamical equations [3]. For instance, the equation for a is

$$a = \sum_j R_j \ast \frac{d^2 s_j}{dt^2} - \sum_j \omega_j^2 \ast s_j + \sum_j A_j \ast g_j^c + g + a_r. \quad (2)$$

The symbol \ast indicates time convolution. R_j in Eq. (2) is a linear operator that represents the unwanted pickup, by the differential interferometer, of generalized coordinates other than x_{12} , like for instance x_1 . These coordinates are measured by the signals s_j , just as s_{12} measures the coordinate x_{12} . Ideally R_j should be zero, but imperfections and misalignments make it nonzero.

In principle R_j acts on coordinates, not on signals. Substituting coordinates with signals, produces an extra term, as signals are always affected by some readout noise. We absorb this term into the overall readout noise a_r .

The readout noise, because of the second time derivation, raises, in power, with the frequency f as $\sim f^4$, and is expected to dominate the data above some 30 mHz, thus setting LPF's measurement bandwidth.

The generalized differential forces *per unit mass*, appearing on the right-hand side of Eq. (2) are split into three contributions:

- (i) The linear operator ω_j^2 converts the relative motion of the TMs and the satellite, along any of the degrees of freedom, into a differential force along x . For really free-falling TMs, ω_j^2 should be zero. However static force gradients within the satellite makes the diagonal coefficients nonzero, while various kind of imperfections and misalignments contribute to nondiagonal terms.
- (ii) The forces commanded by the control loops g_j^c , that are converted into true forces by the linear “calibration” operator A_j . A_j should be just 1 when j corresponds to the electrostatic force commanded along x on TM2, and zero for any other value of j , but deviates from that because of imperfect calibration, delays and signal cross-talk.
- (iii) The random forces g stemming from all remaining disturbances, and whose measurement is the primary target of LPF.

Measuring R_j and A_j is one of the tasks of our analysis. Furthermore, as the coupling of the TMs to the satellite is expected to be present also in a GW observatory like eLISA [4], one of the goals of LPF is to give a measurement of ω_j^2 to be compared with the prediction of the physical model of the system.

The most important goal for LPF though, is that of measuring the PSD of g , the parasitic forces that act on TMs and push them away from their geodesic trajectories.

Equation (2) suggests a natural way of achieving both these goals. Indeed both the s_j ’s and the g_j^c ’s are known, as the first have been measured, and the second have been commanded by the control loops. Thus a fit of the s_j ’s and the g_j^c ’s to a , returns R_j , ω_j^2 , and A_j , as best fit parameters, but also allows the estimation of the PSD of g from the fit residuals, that is from the difference between the acceleration data series and the fitting model.

In reality we need to perform such fits on the data from two different kinds of experiment.

When the target is that of measuring, with comparatively high precision, the values of R_j , ω_j^2 and A_j [see Eq. (2)], we perform dedicated calibration campaigns, where some proper guidance signals are injected into the appropriate control loops, so that the s_j ’s and the g_j^c ’s undergo large variations. This way R_j , ω_j^2 and A_j can be measured with large signal-to-noise Ratio (SNR).

When the target is instead a higher accuracy measurement of the PSD of the ultimate background acceleration noise, we do not apply any guidance signal, but just record acceleration noise data. These data are then fit to the s_j ’s and the g_j^c ’s, with the aim of separating g from the effect of the other force terms in the right-hand side of Eq. (2). Indeed g becomes now the residual of the fit, that is, the difference between the acceleration data and the best fit

model. Actually, also other time series, like thermometer or magnetometer data, may be fitted to the acceleration data to detect and separate specific disturbance sources.

It is worth stressing that, the s_j ’s and the g_j^c ’s cannot be turned off at any time so that an independent measurement of g cannot be performed. A similar situation would also hold for a GW detector like eLISA, where large signals are expected to dominate the data at all times so that an independent measurement of the background noise cannot be performed.

To perform these fits we could not use a standard least-squares method, and we had to develop a different fitting method. Indeed, to perform a least-squares fit on data with colored background noise, as is certainly the case for LPF, one needs an *a priori* knowledge of the background noise PSD, either to set up a whitening filter, if the fit is performed in the time domain, or, for the more common case of a fit in the frequency domain, to assign the statistical weights to each fit residual. However, in our case, the PSD is not known *a priori* and is actually one of the targeted outputs of the fit.

We have then developed a fitting method that works without an *a priori* knowledge of the background noise PSD. The method returns, besides the value of the fitting parameters, also an estimate for the background noise PSD. To achieve a comparatively high precision PSD estimation, the method preserves the ability of averaging over independent data stretches, like with the standard Welch’s averaged periodogram technique [5]. We use this method in the framework of Bayesian estimation. However, we show in the paper that it can also be extended to the standard “frequentist” fitting approach.

Over the last few years, different authors, in the framework of GW detection and Bayesian parameter estimation, have addressed the problem of fitting without a complete *a priori* knowledge of the noise PSD [6–8]. The emphasis of these studies was mostly on minimising the bias that such a lack of knowledge may induce in the estimated signal parameters. This is a different target than the one we are discussing here, where the estimation of the noise is the main goal of the measurement, but the essence of the problem is the same.

Two main approaches have been followed in these studies:

- (1) Within the first approach [6,9], the value of the noise PSD $S_k \equiv S(f_k)$, at each discrete frequency $f_k \equiv k/NT$, with N the length of the data series and T the sampling time, is assumed to be described by some relatively smooth function of frequency, also depending on a vector of some adjustable parameters $\vec{\eta}$, $S_k = S(f_k, \vec{\eta})$.

The likelihood of the fit residuals becomes then a function both of signal parameters and of $\vec{\eta}$.

Appropriate prior probability densities—often some broad Gaussian or uniform densities—are then chosen for both the signal parameters and $\vec{\eta}$. Finally

the posterior likelihood for all parameters is numerically derived by the Markov chain Monte Carlo (MCMC) technique.

Once the global likelihood has been derived, the marginal likelihood of the signal parameters alone, can be derived by numerically marginalizing over the $\vec{\eta}$.

- (2) In the second approach [7,8] the values of the PSD at each discrete frequency S_k , are considered as independent parameters of the likelihood, each one distributed with a prior in the form of a scaled inverse χ^2 density, a family of distributions describing the statistics of the reciprocal of the square of Gaussian variables, that depend on two characteristic parameters [8,10]. This way the posterior density of both the signal parameters and of the S_k , becomes an analytical function of the observed residuals and of the prior parameters. Once the values for these prior parameters have been chosen for each frequency, and the authors discuss possible criteria for this choice, the likelihood can be calculated numerically by MCMC.

Our approach is close to the one in point 2 above with the following main differences and/or extensions:

- (1) We adopt, for the S_k , a family of priors that are uniform, either in the logarithm or in some small power of S_k , over a wide, but finite range of values. These priors give a realistic representation of our knowledge on the residual noise of the system (see sect. II B). The infinite range counterpart of these uniform priors can be obtained from the scaled inverse χ^2 family for particular values of the prior parameters.
- (2) With this assumption we are able to extend the method to the very important case where the time domain data are partitioned into (overlapping) stretches, so that the standard Welch's averaged, and windowed, periodogram of the residuals can be used for the fit [5].

We show that, by using this approach,

- (a) The posterior likelihood can be *analytically* marginalized over the S_k 's so that the marginalized likelihood of signal parameters, which takes a very simple form, can be easily calculated numerically by MCMC, or numerically maximized within a standard fitting approach.
- (b) S_k can then be estimated analytically, and this estimate is shown to be consistent, its error tending to zero as the inverse square root of the number of averaged periodograms.
- (c) The above estimate shows a slight bias that depends on the specific prior adopted, but this bias tends to zero linearly with the inverse of the number of averaged periodograms.

The paper presents such a method and is organized as follows. In Sec. II we describe the method. In Sec. III we give a test of the method with synthetic data of known PSD, and we present a few examples of its application to the reduction of data from LPF end-to-end mission simulator. Finally, in Sec. IV we briefly discuss the results and the possibility of extending the method to signal extraction for the data of GW detectors.

II. MAXIMUM LIKELIHOOD WITH UNKNOWN COLORED NOISE

Though the method is general, for the sake of clarity we will continue to refer to the example of LPF. The main signal for LPF is the relative acceleration data series $a[n]$. We assume that the acceleration data series may be modeled as

$$a[n] = g_m[n, \vec{\theta}] + g[n]. \quad (3)$$

In Eq. (3) the data series $g_m[n, \vec{\theta}]$ consists of the samples of a linear combinations of measured signals, like for example $\omega_j^2 s_j$ or $A_j g_j^c$ in Eq. (2). $g_m[n, \vec{\theta}]$ may have been obtained by processing those signals by some algorithm that depends on a set of parameters θ_i that we have arranged into the vector $\vec{\theta}$. For instance, the force commanded on TM2 along x is described by $A g_2^c(t - \tau_2)$, where $g_2^c(t)$ is the recorded commanded force. Thus $\vec{\theta}$ contains in this case the calibration amplitude A (nominally one) and the delay τ_2 , with the effect of the latter calculated by numerical interpolation of the data.

$g[n]$ represents the residual differential force noise time series, the main objective of the measurement, though corrupted by the superposition of the time series of the readout noise $a_r[n]$. As discussed before, Eq. (3) indicates that $g[n]$ might be derived from $g[n, \vec{\theta}] \equiv a[n] - g_m[n, \vec{\theta}]$, the time series of "residuals," a series that depends parametrically on $\vec{\theta}$.

The equality in Eq. (3) is preserved by moving to the Fourier domain so that one can write

$$\tilde{g}[k, \vec{\theta}] = \tilde{a}[k] - \tilde{g}_m[k, \vec{\theta}], \quad (4)$$

where the tilde indicates a discrete Fourier transform (DFT).

We define the DFT of a stretch of N data of any series $y[n]$ as

$$\tilde{y}[k] = \frac{1}{\sqrt{N}} \sum_{n=0}^{N-1} y[n] w[n] e^{-i n k \frac{2\pi}{N}}. \quad (5)$$

In the transformation we have already included the multiplication of data by a properly selected spectral window $w[n]$. This is common practice in spectral estimation to avoid excess spectral leakage [11].

A. Building up the likelihood function

We now want to discuss the joint probability density function of the residuals $\tilde{g}[k, \vec{\theta}]$, conditional to a specific choice of the values of the θ 's. We will assume here that all noise sources are Gaussian and zero mean, so that also the $\tilde{g}[k, \vec{\theta}]$ are zero-mean and Gaussian. Large non-Gaussian noise in LPF, like glitches and spikes, can satisfactorily be treated as being constituted of signals and subtracted from the data with minimal corruption of the data, given the focus on the lowest frequencies.

The residuals $\tilde{g}[k]$ are complex quantities. It is shown in Appendix A that, for $|k| \geq k_o$ and $|k - k'| \geq k_1$, where k_o and k_1 are integers depending on the adopted spectral window, $\text{Re}\{\tilde{g}[k]\}$, $\text{Im}\{\tilde{g}[k]\}$, $\text{Re}\{\tilde{g}[k']\}$, and $\text{Im}\{\tilde{g}[k']\}$ may be considered, with good accuracy, as all zero-mean, independent Gaussian variables. For instance, for the Blackman-Harris spectral window we commonly use in LPF data analysis, we show in Appendix A, that a safe assumption is $k_o = k_1 = 4$. The variances of $\text{Re}\{\tilde{g}[k]\}$, and of $\text{Im}\{\tilde{g}[k]\}$, are given by (see Appendix A)

$$\sigma_{\text{Re}\{\tilde{g}[k]\}}^2 = \sigma_{\text{Im}\{\tilde{g}[k]\}}^2 = \frac{1}{2} S_k, \quad (6)$$

where S_k is the frequency averaged discrete time PSD of $g[n]$, at the frequency $f_k = k/NT$ defined as

$$S_k = \frac{1}{2\pi} \int_{-\pi}^{\pi} \tilde{S}_{\tilde{g}}(\phi) \left| w\left(\phi - k \frac{2\pi}{N}\right) \right|^2 d\phi. \quad (7)$$

Here $w(\phi)$ is the discrete-time Fourier transform of $w[n]$ and $\tilde{S}_{\tilde{g}}(\phi)$ is the discrete-time power spectral density of the *infinite length* $g[n]$ series, from which the set of N data under analysis has been extracted. Notice that, if aliasing is avoided, then $\tilde{S}_{\tilde{g}}(\phi) = TS(f = \phi/2\pi T)$, where $S(f)$ is the ordinary PSD of the continuous process $g(t)$ of which $g[n] = g(t = nT)$ constitute the series of the samples.

Under the hypotheses above, the joint conditional probability density function of $\text{Re}\{\tilde{g}[k]\}$ and $\text{Im}\{\tilde{g}[k]\}$ is given by

$$P(\vec{g}|\vec{\theta}, \vec{S}) = \prod_{k \in Q} \frac{1}{\pi S_k} e^{-\frac{|\tilde{g}[k, \vec{\theta}]|^2}{S_k}}. \quad (8)$$

In Eq. (8), Q is the subset $Q = \{k_0, k_0 + k_1, k_0 + 2k_1, \dots\}$, of the integer set $0 \leq k \leq N/2$. In the same equation, and in the rest of the paper, we have organized the residuals $\tilde{g}[k, \vec{\theta}]$ and the S_k , with $k \in Q$, into the vectors \vec{g} and \vec{S} , respectively.

It is standard, in spectral analysis, to partition data series into shorter stretches, and to average the spectral estimate over these stretches. Different stretches are treated as statistically independent, even in the presence of partial overlap between adjoining stretches, if these have been tapered at their ends with a proper spectral window.

Assuming that the data series have been partitioned into N_s independent stretches, the probability density in Eq. (8) becomes

$$P(\vec{g}|\vec{\theta}, \vec{S}) = \prod_{k \in Q} \frac{1}{(\pi S_k)^{N_s}} e^{-N_s \frac{|\tilde{g}[k, \vec{\theta}]|^2}{S_k}}, \quad (9)$$

where the bar represents an average over the N_s stretches.

In Eq. (9) we have written the probability density of the data as also being conditional on \vec{S} . In a standard fitting procedure, these coefficients are assumed to be known. On the contrary, here we discuss the case where the components of \vec{S} are unknown and must be estimated from the data.

We now write down the posterior probability density of \vec{S} and $\vec{\theta}$,

$$P(\vec{\theta}, \vec{S}|\vec{g}) = \frac{P(\vec{g}|\vec{\theta}, \vec{S}) \times \prod_{k \in Q} P(S_k) \times P(\vec{\theta})}{\int P(\vec{g}|\vec{\theta}, \vec{S}) \times \prod_{k \in Q} P(S_k) dS_k \times P(\vec{\theta}) d\vec{\theta}}. \quad (10)$$

In Eq. (10) we have made the key assumptions that S_i is independent of S_j if $i \neq j$, and that both are independent of $\vec{\theta}$. This way the joint *prior* probability density $P(\vec{\theta}, \vec{S})$ splits into the product of the separated prior probability densities $P(\vec{\theta})$ and $P(S_k)$.

While the independence of $\vec{\theta}$ and \vec{S} is rather natural, the physical basis for the independence of S_i and S_j , when $i \neq j$, may need some justification.

Our *a priori* knowledge of the noise PSD, in the case of an instrument where the signal cannot be ‘turned off’ and the noise independently measured, is rather limited. Unexpected lines could be present in the spectrum, such that nearby coefficients, S_i and $S_{i \pm 1}$, may differ even in order of magnitude. Thus even the perfect knowledge of S_i would not give us any significant information on the probability density function of any of the other S_j 's, which is the very definition of independent random variables.

B. Fitting

We now discuss the use of the likelihood function in Eq. (10) for the purpose of fitting. As most of our fits involve nonlinear functions of parameters, our preferred approach is that of Bayesian parameter estimation with the MCMC method, but we will also consider a more conventional approach wherein one searches for a maximum of the likelihood as a function of fitting parameters. We will show now that, whatever the selected approach, with a proper choice of the prior probability density of the components of \vec{S} , the parameter space can be reduced to just that of $\vec{\theta}$.

In order to do this, we assume that the prior of S_k is uniform as a function of either some small power of S_k , or of $\log(S_k)$, between two values $S_{k,a} \ll S_{k,b}$. This is not

exactly the same as using Jeffrey's noninformative prior [12], though it may be approximated by this under some assumptions that we will make later.

Such a choice for the prior, in particular a uniform density as a function of $\log(S_k)$, again closely reflects our *a priori* physical knowledge of the noise. Indeed, as said, in LPF, as in a GW observatory like eLISA, signals cannot be turned off and the background noise cannot be independently measured. Though we may have theoretical models for the sources of g , unexpected noise sources may lead to a PSD deviating from theoretical expectations even in order of magnitude.

With this choice of the prior, $P(\vec{\theta}, \vec{S}|\vec{g})$ can be analytically integrated over the \vec{S} space, to obtain a marginalized likelihood which is only a function of $\vec{\theta}$. By assuming that $P(\vec{\theta})$ is bound to a domain on which $|\tilde{g}[k]|^2 \ll S_{k,b}$, and that, for any value of $\vec{\theta}$, $S_{k,a} \ll |\tilde{g}[k]|^2$, the integration over S_k can be extended from zero to infinity:

$$P_{\text{marg}}(\vec{\theta}|\vec{g}) = \int_0^\infty P(\vec{\theta}, \vec{S}|\vec{g}) d\vec{S}. \quad (11)$$

By performing the integration we obtain

$$P_{\text{marg}}(\vec{\theta}|\vec{g}) = \frac{\prod_{k \in Q} \overline{|\tilde{g}[k, \vec{\theta}]|^2}^{m-N_s}}{\int \prod_{j \in Q} \overline{|\tilde{g}[j, \vec{\theta}']|^2}^{m-N_s} P(\vec{\theta}') d\vec{\theta}'}, \quad (12)$$

for a uniform prior either in S_k^m or in $\log(S_k)$, in which case one should put $m = 0$ in the formulas.

In addition, the likelihood $P(\vec{\theta}, \vec{S}|\vec{g})$, for any given value of $\vec{\theta}$, reaches a maximum $P_{\text{max}}(\vec{\theta}|\vec{g})$ for some values $S_{k,\text{max}}$ of the S_k 's. Both the value of $P_{\text{max}}(\vec{\theta}|\vec{g})$, and of $S_{k,\text{max}}$ can be calculated analytically by differentiation of Eq. (9). We get

$$S_{k,\text{max}} = \frac{N_s}{N_s - m + 1} \overline{|\tilde{g}[k, \vec{\theta}]|^2}, \quad (13)$$

and

$$P_{\text{max}}(\vec{\theta}|\vec{g}) = \frac{e^{-(N_s+1)} (N_s + 1)^{N_s+1-m}}{N_s (N_s - m - 1)!} \frac{P_{\text{marg}}(\vec{\theta}|\vec{g})}{\prod_{k \in Q} \overline{|\tilde{g}[k, \vec{\theta}]|^2}}. \quad (14)$$

We note that Eq. (13), in the limit of large N_s , where all priors give the same formula, anticipates our result for the estimate of S_k that we further discuss in detail in subsection II D.

These results allow us to restrict the fitting to the $\vec{\theta}$ parameter space. In the maximization approach, one can search for a maximum of $P_{\text{max}}(\vec{\theta}|\vec{g})$, that will also be a

maximum of $P(\vec{\theta}, \vec{S}|\vec{g})$. Within the Bayesian approach, the likelihood mapping may be performed over just $\vec{\theta}$, by using $P_{\text{marg}}(\vec{\theta}|\vec{g})$.

In both cases, the results in Eqs. (12), (13), and (14) indicate that, for large enough N_s the logarithm of the likelihood to be either maximized, or used in MCMC mapping, is

$$\Lambda(\vec{\theta}) \equiv \log(P_{\text{marg}}(\vec{\theta}|\vec{g})) = -N_s \sum_{k \in Q} \log(\overline{|\tilde{g}[k, \vec{\theta}]|^2}) + C, \quad (15)$$

instead of the standard least-squares fitting result with known S_k ,

$$\Lambda(\vec{\theta}) = -N_s \sum_{k \in Q} \frac{\overline{|\tilde{g}[k, \vec{\theta}]|^2}}{S_k} + C'. \quad (16)$$

Here, C and C' are just constants. In essence, according to Eq. (15), in the presence of unknown and unmodeled noise, any fit must minimize not the mean square residuals, but rather the sum of their logarithm. This is one of the main results of the paper.

As already mentioned in the introduction, a likelihood proportional to that in Eq. (12) has been found for the special case $N_s = 1$ and $m = 0$ by [8]. Our result generalizes it to the experimentally important case of averaged and windowed periodograms, and to the case of $m \leq 0$.

For the rest of the paper we will call the likelihood in Eq. (15), the “logarithmic” likelihood (LL).

C. Truly independent DFT coefficients and spectral resolution

In order to maintain the accuracy of the result in Eq. (15), one should fulfil the condition $|k - k'| \geq k_1$, dropping a great number of DFT coefficients, and thus lowering the spectral resolution of the fit. This might be undesirable at the lowest frequencies, where the relative spectral resolution is rather low. The inaccuracy deriving from summing over *all* DFT coefficients in Eq. (15) reduces to counting each independent DFT coefficient more than once. Thus, using such a likelihood function may overestimate the number of degrees of freedom of the fit and may lead to an underestimate of the parameter errors. This effect may be corrected for by scaling down the likelihood by an appropriate correction factor γ , that is by using the likelihood $\tilde{\Lambda}(\vec{\theta}) \equiv \gamma \times \Lambda(\vec{\theta})$. This is further discussed in Sec. III.

As for the coefficients with $k < k_0$, in spectral estimation these are discarded in any case, because of the strong bias they suffer from the leakage of spectral power from frequencies below $f = 1/NT$.

D. Summary of fitting procedure and the uncertainty on parameter estimates

The derivation in the preceding sections leads to a very simple implementation of the method. The entire machinery of a standard frequency domain fit can be retained, provided that the usual likelihood function is replaced with that in Eq. (15).

Thus, for instance, within the framework of Bayesian estimation, rather than sampling from the joint likelihood of $\vec{\theta}$ and \vec{S} , parameter estimates of $\vec{\theta}$ are obtained from the marginalized LL likelihood, which can be numerically mapped over the space of the θ 's by standard MCMC.

Such a calculation does not return an estimate for the S_k and their errors. These however can be estimated as follows. Let's write the conditional mean value of any integer power of S_j , S_j^m as

$$\langle S_j^m | \vec{g} \rangle = \frac{\int P(\vec{\theta}) d\vec{\theta} \int_0^\infty S_j^m P(\vec{g} | \vec{\theta}, \vec{S}) P(\vec{S}) d\vec{S}}{\int P(\vec{\theta}) d\vec{\theta} \int_0^\infty P(\vec{g} | \vec{\theta}, \vec{S}) P(\vec{S}) d\vec{S}}. \quad (17)$$

This integral may be somewhat reduced if one uses the prior for \vec{S} discussed above. Indeed one can check (see Appendix B) that

$$\langle S_j | \vec{g} \rangle = \frac{N_s}{N_s - m - 1} \langle |\vec{g}[j]|^2 \rangle \quad (18)$$

and that

$$\langle S_j^2 | \vec{g} \rangle = \frac{N_s^2}{(N_s - m - 1)(N_s - m - 2)} \langle |\vec{g}[j]|^4 \rangle. \quad (19)$$

Here the mean value $\langle |\vec{g}[j]|^4 \rangle$ is the value of $|\vec{g}[j, \vec{\theta}]|^4$ averaged over the total posterior probability density of $\vec{\theta}$. Thus, once this posterior probability has been calculated as described above, in principle also the values of S_j and of S_j^2 , and then of the variance of S_j , can also be numerically calculated.

The numerical calculation can be avoided if the parameter posterior distribution is symmetric enough around the mean values, a condition we always find satisfied in our numerical calculations. In this case the mean values in the rightmost terms of the Eq. (18) and Eq. (19) will coincide with the values they take at the mean values $\vec{\theta} = \vec{\theta}_o$ where the likelihood is also a maximum. Then, for $N_s \gg 1$, m :

$$\begin{aligned} \langle S_j | \vec{g} \rangle &\approx |\vec{g}[j, \vec{\theta}_o]|^2 \\ \sigma_{S_j} &\equiv \sqrt{\langle S_j^2 | \vec{g} \rangle - \langle S_j | \vec{g} \rangle^2} \approx \frac{|\vec{g}[j, \vec{\theta}_o]|^2}{\sqrt{N_s}}. \end{aligned} \quad (20)$$

Equation (20) shows that the estimate of the S_j is consistent, as its error decreases as $1/\sqrt{N_s}$.

It is worth noticing that, by increasing N_s , besides improving the precision of the noise estimate, one also rapidly suppresses any residual dependance of both the likelihood in Eq. (12) and the momenta in Eqs. (18) and (19), on m , the only remaining parameter in our approach that is still needed to characterize the noise. Indeed marginalization removes the dependance of the likelihood on the S_k 's but does not remove its dependence on parameters, like m , entering in the prior for S_k , whose value must then be decided in advance (see for instance Ref. [8]). It appears then that averaging over many periodograms may be an effective way to remove such a residual bias.

In this sense, neither the estimate for S_k nor that for $\vec{\theta}$ are biased by noise modeling.

Thus, in summary, the fitting proceeds in two steps: after having mapped the marginalized likelihood LL to estimate $\vec{\theta}$, \vec{S} and their uncertainties can be evaluated, through Eq. (20), from the fit residuals evaluated for $\vec{\theta} = \vec{\theta}_o$, where LL takes its maximum value.

E. Relation to the iteratively reweighted least-squares (IRLS) method

A popular method to perform least-square fitting without an *a priori* knowledge of the expected residual PSD, is that of performing the fit iteratively. The procedure starts by performing an ordinary least-squares fit, maximizing then the likelihood in Eq. (16), by using some arbitrary initial value for the S_k , like for instance $S_k = 1$, or alternatively by using the spectrum of the raw data before fitting. Once the maximum has been found at some point $\vec{\theta}_o$, the square modulus of the residuals of this first fit, $|\vec{g}[k, \vec{\theta}_o]|^2$ are used as the S_k 's in a new run of the fit. Often the residuals are smoothed over the frequency, or fitted to some smooth model before using them as weights in the next iteration of the fit. However the use of the residuals as they are is unbiased and numerically lighter, thus we prefer to restrain the discussion to just this case. The procedure is then iterated until the residuals of the fit and the S_k do not change anymore to within some tolerance. The procedure usually converges quite rapidly. We will show here that the parameter values that are obtained with this iterative procedure are the same as those obtained by a LL.

Consider that, at the n^{th} step of the procedure, one wants to minimize

$$\chi_n^2 \equiv N_s \sum_{k \in Q} \frac{|\vec{g}[k, \vec{\theta}_n]|^2}{|\vec{g}[k, \vec{\theta}_{n-1}]|^2} \quad (21)$$

as function of $\vec{\theta}$. If the method converges, at some point the sequence becomes stationary and independent of n . Let us write

$$\overline{|\tilde{g}[k, \vec{\theta}_n]|^2} \approx \overline{|\tilde{g}[k, \vec{\theta}_{n-1}]|^2} + \sum_i \frac{\partial \overline{|\tilde{g}[k, \vec{\theta}_{n-1}]|^2}}{\partial \theta_i} \delta \theta_i, \quad (22)$$

and substituting in Eq. (21), we get

$$\begin{aligned} \chi_n^2 &\approx N_s \sum_{k \in Q} \left(1 + \sum_i \frac{\partial \overline{|\tilde{g}[k, \vec{\theta}_{n-1}]|^2} / \partial \theta_i}{\overline{|\tilde{g}[k, \vec{\theta}_{n-1}]|^2}} \delta \theta_i \right) \\ &= N_Q N_s + \sum_i \frac{\partial N_s \sum_{k \in Q} \log(\overline{|\tilde{g}[k, \vec{\theta}_{n-1}]|^2})}{\partial \theta_i} \delta \theta_i, \end{aligned} \quad (23)$$

where N_Q is the number of elements in Q . If the sequence χ_n^2 must become stationary, then from some value of n on, the derivatives in the right-hand side of Eq. (23) must become zero. Then $N_s \sum_{k \in Q} \log(\overline{|\tilde{g}[k, \vec{\theta}_{n-1}]|^2})$ reaches a maximum and hence so does the likelihood in Eq. (15).

This shows that IRLS and LL achieve the same estimate for the best fit parameters. This will be also shown numerically in Sec. III.

For nonlinear fits, when the maximum of the likelihood must be searched for numerically, the straightforward maximization of the likelihood in Eq. (15) is substantially quicker than the IRLS method, as the maximization with respect to the parameters must be performed only once. On the other hand, when the dependence of the likelihood on the parameters is linear, that is if the θ 's are just amplitude coefficients, then the maximization, at each step, reduces to solving a system of linear equations and the IRLS method may be substantially faster.

A theoretical estimate of parameter errors with the IRLS method, is not straightforward. For nonlinear fit models, once the stationary solution has been reached, one can use the residuals from this solution in place of the S_k in the least-squares likelihood of Eq. (16) and perform an MCMC map, from which the parameter errors can be calculated. If the fit model is linear, one can apply the standard error propagation, or Fisher matrix technique, to the stationary solution and derive the errors from that. Though this is an heuristic approach, we will show in Sec. III, that, at least in the case of the nonlinear models that we have studied in the present paper, it gives results consistent with those obtained with the LL fits.

Finally we want to make a couple of remarks:

- (i) The equivalence of the two methods shows that an IRLS fit that uses the single DFT coefficients from one fit to weight the next fit—as opposed to taking some smooth, few parameter model of the spectrum frequency dependence, such as a polynomial—implicitly assumes that the S_k are uncorrelated and uniformly distributed in some logarithmic or power law space, as our LL method explicitly does.
- (ii) The assumption of uncorrelated and uniformly distributed S_k allows for arbitrary variations of the

PSD from one frequency to the next one and allows us then to weight neighboring frequencies quite differently based on the DFT of their residuals. The impact of this inhomogeneous weighting is limited in practice by averaging the spectrum over many time windows, which reduces the fluctuations in the mean square residuals between neighboring bins, and by using a small number of disturbance parameters θ . While we see no evidence of anomalies in the results of our analyses, we bear in mind that any assumption about the S_k can have an impact on the analysis. Our technique allows for an inhomogeneous weighting of points, as is consistent with our hypothesis of uncorrelated and uniformly distributed S_k , and, unlike a low-order polynomial model of the spectrum frequency dependence, allows for lines, bumps, and other unexpected spectral features.

III. APPLICATION TO DATA

We have tested the method in two ways:

- (i) With the aim of verifying its accuracy, we have performed a noise decomposition exercise on a set of data generated from a fully known PSD.
- (ii) In order to test the practical usability of the method with nonideal data, we have used it in various data reduction exercises on the data from the LPF end-to-end mission simulator.

In the next sections we describe these tests and report on their results, while in Sec. IV we discuss their significance.

A. Test on data with a known PSD

To perform this test we have generated some data series with spectral content and dynamic range similar to those of LPF data. A base data series has been created which is the sum of two components. The first, $a[n]$, simulating the acceleration data, is the numerical second derivative, obtained as described in [2], of a “red” noise series. This red series is composed of a $1/f^6$ PSD low frequency tail merging, at 10 mHz, into a flat PSD. The second series, $g^c[n]$, which simulates the feedback force data, consists of a tail with $1/f^4$ PSD merging into a flat plateau at 0.3 mHz. This base series is contaminated by independently superposing a series $\alpha_{1,o} \delta a[n]$ to $a[n]$ before derivation and one $\alpha_{2,o} \delta g[n]$ to $g^c[n]$. $\delta a[n]$ has a $1/f^6$ PSD, while $\delta g[n]$ has a band-pass structure with a broad peak around 2 mHz. Thus the data consist of

$$g[n] = a[n] + g^c[n] + \alpha_{1,o} \frac{d^2 \delta a[n]}{dt^2} + \alpha_{2,o} \delta g[n]. \quad (24)$$

We have then prepared two data series to be used for the fit and the noise decomposition.




FIG. 2 (color online). Result of noise projection with simulated data. Upper noisy red line: the square root PSD of the base data series $g[n]$. Lower noisy black line: the square root PSD of the residuals after fitting $g[n]$, with $\alpha_1 \frac{d^2 \delta a}{dt^2}[n] + \alpha_2 \delta g[n, \tau]$. The solid blue line represents the theoretical value of the expected PSD.

The first is the double derivative $\frac{d^2 \delta a}{dt^2}[n]$ of $\delta a[n]$. The second is a copy $\delta g[n, \tau]$ of $\delta g[n]$ but delayed by an amount τ .

A fit to $g[n]$ was then performed in the frequency domain with the fitting function

$$g^{\text{fit}}[n] = \alpha_1 \frac{d^2 \delta a}{dt^2}[n] + \alpha_2 \delta g[n, \tau], \quad (25)$$

with α_1 , α_2 , and τ being the free parameters of the fit.

All data series have been prepared with a sampling time of 10 Hz, and then, after having applied the appropriate delays, low-passed and decimated to 1 Hz. In analogy to the LPF data analysis, we further low-passed and decimated the

data to 0.1 Hz. We only fit the DFT coefficients with frequencies $f \leq 50$ mHz.

Furthermore we have used $N_s = 9$ data stretches with 50% overlap, and a Blackman-Harris spectral window with $k_o = 4$. Finally we have used $k_1 = 1$.

Fig. 2 shows a typical result of the fit.

The fit has been performed with the MCMC method and the Metropolis-Hastings algorithm with the LL. The parameter values used for the plot in Fig. 2 correspond to the maximum of the likelihood for that specific realization of the data. An example of the marginalized MCMC distributions of the different fitting parameters is reported in Fig. 3.

The simulation has been repeated $N_{\text{rep}} = 40$ times. Each time we have generated a new data series to which we have applied the MCMC fit. For each repetition we have recorded the mean values and the standard deviations of the MCMC distributions of the three fitting parameters.

For comparison we have also performed similar simulations for the following cases:

- (i) $k_1 = 4$ instead of $k_1 = 1$.
- (ii) $k_1 = 1$ and $\gamma = 1/2$.
- (iii) $k_1 = 1$ and $\gamma = 1/3$.
- (iv) $k_1 = 1$, but with the IRLS method. In this case, in each simulation, the IRLS fit has been performed with a numerical minimization routine until a stationary solution has been reached. Then an MCMC sequence has been generated by using the likelihood in Eq. (16), with the S_k equal to residuals of the stationary solution.

The results are summarized in Table I.

It is worth pointing out the following facts that can be observed in Table I:

- (i) For all fitting parameters, the average values resulting from the different fitting methods, agree with each other and, with the true values, within their




FIG. 3. Marginalized histograms for the three fitting parameters, obtained with the MCMC method and the LL.

TABLE I. Results from simulations. The meaning of the columns is the following. True: parameter values used in the simulation. Columns under the header “LL” refer to fits performed with the LL method. Columns under the header ‘IRLS’ refer to fits performed with the IRLS method. Av.: average of the 40 mean values of the MCMC distributions obtained in the 40 independent simulations. σ_{sample} : standard deviation of the 40 mean values. σ_{MCMC} : root mean square of the standard deviations of the MCMC distributions of the 40 independent simulations.

Param.	True	LL									IRLS					
		$k_1 = 1, \gamma = 1$			$k_1 = 4, \gamma = 1$			$k_1 = 1, \gamma = 1/2$			$k_1 = 1, \gamma = 1/3$			$k_1 = 1, \gamma = 1$		
		Av.	σ_{sample}	σ_{MCMC}	Av.	σ_{sample}	σ_{MCMC}	Av.	σ_{sample}	σ_{MCMC}	Av.	σ_{sample}	σ_{MCMC}	Av.	σ_{sample}	σ_{MCMC}
α_1	300	300	7.9	5.8	300	12	12	302	10	8.4	301	8.9	10	297	9.8	5.5
α_2	2	2.005	0.026	0.014	2.004	0.029	0.027	2.005	0.021	0.019	2.001	0.026	0.023	2.000	0.025	0.013
$\tau[s]$	0.8	0.84	0.29	0.18	0.70	0.42	0.36	0.79	0.34	0.26	0.77	0.37	0.31	0.90	0.33	0.17

respective errors $\sigma_{\text{sample}}/\sqrt{N_{\text{rep}}}$, with the exception of the results of IRLS simulation. These results however agree with the rest at worst to within $1.8\sigma_{\text{sample}}/\sqrt{N_{\text{rep}}}$. In particular this confirms that the LL method produces an unbiased estimate of the parameters.

- (ii) On the contrary, for $\gamma = 1$, σ_{sample} and σ_{MCMC} agree for all parameters, within their relative uncertainties, *only for the case* $k_1 = 4$. In this last case however, both σ_{sample} and σ_{MCMC} are somewhat larger than those obtained with $k_1 = 1$, both for case of the LL and for that of IRLS.
- (iii) The agreement between σ_{sample} and σ_{MCMC} may be recovered also for $k_1 = 1$, if $\frac{1}{3} < \gamma < \frac{1}{2}$.
- (iv) The IRLS procedure followed by the MCMC likelihood map, takes approximately 4 times the time needed by the straightforward LL, MCMC map.

B. Application to LPF simulator data analysis

As is customary with space missions, an end-to-end mission simulator of LPF has been set up by industry [13]. The simulator includes not only the linear dynamics of satellite and TMs translation, but also realistic, nonlinear models of the critical parts of the system, like the electrostatic actuation system, the rotational dynamics of both the TMs and the satellite, the interferometers, etc. These nonlinearities are not expected to play a significant role, as all displacements, velocities, etc. are expected to be very small during science operations. We expect then to be able to understand the largest part of LPF results within a linear model, and to have to deal with nonlinearities only as occasional small deviations from the linear regime. The simulator is proving extremely useful for testing this approach and the related data analysis algorithms and tools [14].

In this section we illustrate the application of the method described in Sec. II, to some selected cases, taken from the extensive simulation campaigns that have been performed in preparation for the mission operations. Specifically, we first discuss a simulated instrument calibration, where large calibration guidance signals are injected in the proper

control loops. We then give an example of true noise decomposition performed to hunt for the source of some extra noise found in the simulated data.

1. Extraction of calibration signals

As explained in Sec. I, TM2 is forced, by a weak electrostatic force control loop driven by the main interferometer signal s_{12} , to stay nominally at a fixed distance from the reference TM.

This loop compensates most of the low frequency (< 10 mHz) forces g by applying commanded forces g^c . Thus, accurate subtraction of these applied forces, including calibration of the actuator, is needed to extract the disturbance forces. Our analysis subtracts the commanded force series multiplied by a calibration factor, Ag_2^c . A is extracted with a dedicated “system identification” experiment in which a comparatively large modulated guidance signal is added to the measured displacement signal, s_{12} , in the control loop. The injected signal has the effect of modulating the distance between the TMs by exerting forces much larger than those exerted by the loop in the presence of just the background noise.

Our technique to calibrate the transfer function consists of fitting the pre-processed commanded force series $Ag_2^c[n, \theta]$ to $a[n]$. As already mentioned, by ‘pre-processed’ we mean here that the series are filtered via an algorithm that depends on some set of parameters, θ , that become free fitting parameters. Commonly for the $g_2^c[n, \theta]$ series, θ only includes the amplitude, A , and a delay, τ , but single pole filters that simulate the response of the actuators have also been tested in the past [14–17]. In addition to the feedback force removal, we also subtract, from the acceleration data, the forces due to the motion of both TMs within the static force gradient in the satellite [see Eq. (2)]. This is predicted to be dominated by the electric gradient due to various voltages applied between the TMs and their surroundings, and by the gravitational gradient. As already explained, these forces are proportional to the TM displacements relative to the satellite, measured by s_1 and s_{12} . In the calibration experiment these displacements are also large, so that the effect of the respective forces may have




FIG. 4 (color online). Deviation of the best fit parameter obtained at each step of the IRLS procedure, from those obtained with the MCMC mapping of the LL. Deviations are divided by the standard deviation obtained from the MCMC chain of the LL. Colors mark the different parameters. Red circles: A . Blue triangles: $-\omega_2^2$. Black diamonds: $-(\omega_1^2 - \omega_2^2)$. Brown stars: τ .

appreciable SNR. As a consequence, we include in the fitting model two terms proportional to s_1 and to s_{12} , respectively. In conclusion we fit the acceleration data series assuming the following model,

$$\tilde{a}[k] = A\tilde{g}_2^s[k, \tau] - \omega_2^2 \tilde{s}_{12}[k] - (\omega_1^2 - \omega_2^2) \tilde{s}_1[k] + \tilde{g}[k], \quad (26)$$

with the amplitudes A , ω_2^2 , $(\omega_1^2 - \omega_2^2)$, and the delay τ as free fitting parameters.

For all the fits, we have used the Blackman-Harris spectral window, taken $N_s = 5$ and limited the frequency range to $f \leq 50$ mHz.

For the sake of comparison, the fit has been performed with both the LL and the IRLS methods. In this last case, the convergence of the best-fit parameter values to those obtained with the LL method has been verified and the results are reported in Fig. 4.

Furthermore, Fig. 5 shows the PSD of the residuals of the two different fits.

Finally, in Table II we report, for the same fits, the parameter values and their uncertainties.




FIG. 5 (color online). PSD of various data series. The upper, blue, noisy continuous line is the PSD of $a[n]$ before the fit. The lower, black, noisy continuous line, barely visible behind the red dashed line, is the PSD of the residuals of the fit with the LL. The red dashed line is the PSD of the residuals of the IRLS fit.

In the case of the IRLS fit, as before, we have first run the reweighted iteration, performing, at each step, just a numerical maximization of the likelihood in Eq. (16). After obtaining the stationary solution for the S_k , we have inserted these values back in the likelihood of Eq. (16) and performed an MCMC map to estimate the parameter errors reported in II. As expected, the results of the two methods are indistinguishable.

It must be stressed that, unfortunately, we cannot straightforwardly compare the results of Table II with the “true” parameter values used in the simulator. Only the parameter $A = 1.05$ has a fixed and traceable value, while both ω_2^2 and ω_1^2 are heuristic single-parameter approximations of a complicated model where the electrical fields, one of the major sources of force gradient, are dynamically calculated as part of the overall three body dynamics of the TMs and the satellite. This is also the case for τ that results from the combination of various propagation delays within the system. All that said, the values resulting from the fit agree with the expectation, within their respective uncertainties.

TABLE II. Parameter values from the various fits of Fig. 5. “LL” refers to the fit performed with the LL. “IRLS” to the fit done with the IRLS method. “Mean” indicates the mean of the MCMC parameter distribution. σ represents its standard deviation.

Parameter	LL		IRLS	
	Mean	σ	Mean	σ
$-\omega_2^2[s^{-2}]$	2.259×10^{-6}	1×10^{-9}	2.259×10^{-6}	1×10^{-9}
$(\omega_2^2 - \omega_1^2)[s^{-2}]$	7.2×10^{-7}	1.2×10^{-7}	7.1×10^{-7}	1.2×10^{-7}
A	1.04998	1.2×10^{-5}	1.04998	1.2×10^{-5}
$\tau[s]$	-4.002×10^{-1}	3×10^{-4}	-4.002×10^{-1}	3×10^{-4}

2. Decomposition of low-frequency noise

The second case of study deals with a typical case of noise hunting and decomposition. We started by forming a time series for $g[n]$, subtracting the commanded force and stiffness effects as calibrated in the previous section. The PSD of $g[n]$ at frequencies below 1 mHz was found in excess of what was expected from the physical model at the basis of the simulator. However calculating the expected PSD is not straightforward, given the complexity of the simulator model, so that no firm conclusion could be made about this noise being a real property of the system or rather a software artefact.

To help in identifying the source, we performed an extensive campaign of decomposition, fitting all available data series to $g[n]$ in an attempt to identify the source. Not knowing what to expect as a background noise, we were forced to develop the methods we are discussing here. Only three data series were observed to significantly reduce the noise when subtracted from $g[n]$. The first and most important series was the difference ΔF_z between the forces applied by the control system to the two different TMs along the z axis (see Fig. 1 for the definition of the axes). We also found a smaller but significant contribution of ΔF_y , the corresponding difference of force along y . We finally found another small contribution from the torque N_x applied to the inertial reference TM along the x axis. Figure 6 shows the effect of subtraction.

In the simulator model, these forces should leak into x via some linear coupling coefficients, the values of which are known simulator input parameters. We found that the values of the coupling coefficients resulting from the fit are in quantitative agreement with those in the simulator, for ΔF_y and N_x (see Table III). The Table also shows that, on the contrary, the coupling coefficient for ΔF_z is ~ 8 times




FIG. 6. The effect of subtraction of ΔF_z , ΔF_y , and N_x . The dotted line represents the PSD of the data series before subtraction. The solid line represents the PSD of the data series after subtraction, i.e. the series of the fit residuals.

TABLE III. Parameter values from LL noise decomposition. Mean and σ are the mean and standard deviation from the MCMC chain. “Expected” refer to the values in the simulator.

Disturbance	Coupling coefficients		
	Mean	σ	Expected
ΔF_z	-8.03×10^{-3}	6×10^{-5}	1.1×10^{-3}
ΔF_y	1.3×10^{-3}	2×10^{-4}	1.1×10^{-3}
N_x	$7.9 \times 10^{-5} m^{-1}$	$9 \times 10^{-6} m^{-1}$	$7.7 \times 10^{-5} m^{-1}$

larger than the corresponding one used in the simulator and, in addition also has the wrong sign.

The named forces and torques are commanded by a controller in charge of stabilizing the absolute orientation of the satellite. Explaining how this controller works goes beyond the scope of the present paper. For the sake of the discussion it is only useful to mention that the controller is driven by the signals from a set of Autonomous Star Trackers, so that our finding allowed us to trace the problem back to an erroneous coupling of these devices.

IV. DISCUSSION AND CONCLUSIONS

The results of the tests reported in Secs. III A and III B 1 show that

- (i) The PSD of the residuals of the fit is in quantitative agreement with the expected spectrum for the background noise.
- (ii) The LL is well behaved and produces unbiased results when used in a MCMC fit.
- (iii) The LL MCMC fit, and the IRLS fit followed by a MCMC likelihood mapping, give the same results, this last one being substantially slower in the case of nonlinear fitting.
- (iv) The estimate of the parameter errors, from the fit performed with $k_1 = 1$, seems indeed to be moderately affected by the correlation between nearby DFT coefficients. This bias, as expected, is common both to LL and IRLS fitting.
- (v) The bias can indeed be made negligible by taking properly spaced coefficients, or by correcting the likelihood with the proper factor γ .

We believe then that the LL fitting presented here is of general use and solves the problem of the lack of *a priori* knowledge of the target noise in frequency domain fits. From the point of view of the computational load, for multiparameter nonlinear fitting, the method is faster than IRLS. However the method is intrinsically nonlinear, and fitting would not reduce to a set of algebraic equations, as IRLS does, when the fitting parameters are just multiplicative amplitudes. Given that the final results of the two method coincide, the IRLS method is preferable in the linear model case.

We think that the LL approach is numerically lighter than the methods [6,18] that employ a parametric model for the

noise in the likelihood in Eq. (8), and then sample the likelihood with MCMC over a parameter space that also includes the noise model parameters. Indeed, with that method the likelihood function contains more terms, the sum of the squares and the sum of the logarithm of the S_k , and the parameter space to be searched numerically is wider.

We think that the method used here, including the partitioning of the data in stretches and the average over the stretches, could be usefully extended also to case of signal extraction from GW detectors. This is particularly true in the case of a space borne detector like eLISA, which is expected to be signal dominated at all times, so that a direct measurement of the instrument noise is difficult. An unbiased estimator of the noise could then help in avoiding the introduction of unwanted bias in the signal parameter estimation.

ACKNOWLEDGMENTS

This work has been supported in part under contracts from Agenzia Spaziale Italiana, Istituto Nazionale di Fisica Nucleare, Deutsches Zentrum für Luft- und Raumfahrt e.V., Swiss Space Office, State Secretariat for Education, Research and Innovation, United Kingdom Space Agency, Plan Nacional del Espacio of the Spanish Ministry of Science and Innovation, and Centre Nationale d'Études Spatiales.

APPENDIX A: PROPERTIES OF DFT COEFFICIENTS

It is a well-known result [11] that

$$\langle |\tilde{g}[k]|^2 \rangle = S_k. \quad (\text{A1})$$

With the same kind of calculations that lead to Eq. (A1), one can easily calculate that

$$\begin{aligned} \langle \text{Re}\{\tilde{g}[k]\}^2 \rangle &= \frac{1}{2\pi} \int_{-\pi}^{\pi} d\phi \tilde{S}_{\tilde{g}}(\phi) \frac{1}{4} \\ &\times \left(\left| w\left(\phi - k\frac{2\pi}{N}\right) \right|^2 + \left| w\left(\phi + k\frac{2\pi}{N}\right) \right|^2 \right. \\ &\left. + 2\text{Re}\left\{ w\left(\phi + k\frac{2\pi}{N}\right) w^*\left(\phi - k\frac{2\pi}{N}\right) \right\} \right), \end{aligned} \quad (\text{A2})$$

$$\begin{aligned} \langle \text{Im}\{\tilde{g}[k]\}^2 \rangle &= \frac{1}{2\pi} \int_{-\pi}^{\pi} d\phi \tilde{S}_{\tilde{g}}(\phi) \frac{1}{4} \\ &\times \left(\left| w\left(\phi - k\frac{2\pi}{N}\right) \right|^2 + \left| w\left(\phi + k\frac{2\pi}{N}\right) \right|^2 \right. \\ &\left. - 2\text{Re}\left\{ w\left(\phi + k\frac{2\pi}{N}\right) w^*\left(\phi - k\frac{2\pi}{N}\right) \right\} \right) d\phi. \end{aligned} \quad (\text{A3})$$

The consequence of Eqs. (A2) and (A3) is that if k is large enough so that $w(\phi - k\frac{2\pi}{N})$ has no overlap with $w(\phi + k\frac{2\pi}{N})$, then

$$\langle \text{Re}\{\tilde{g}[k]\}^2 \rangle = \langle \text{Im}\{\tilde{g}[k]\}^2 \rangle = \frac{1}{2} \langle |\tilde{g}[k]|^2 \rangle = \frac{1}{2} S_k. \quad (\text{A4})$$

In addition one can calculate that

$$\begin{aligned} \langle \text{Re}\{\tilde{g}[k]\} \text{Im}\{\tilde{g}[k]\} \rangle &= \frac{1}{2\pi} \int_{-\pi}^{\pi} d\phi \tilde{S}_{\tilde{g}}(\phi) \frac{1}{2} \text{Im}\left\{ w\left(\phi + k\frac{2\pi}{N}\right) w^*\left(\phi - k\frac{2\pi}{N}\right) \right\}. \end{aligned} \quad (\text{A5})$$

So that again, if k is large enough, then

$$\langle \text{Re}\{\tilde{g}[k]\} \text{Im}\{\tilde{g}[k]\} \rangle = 0. \quad (\text{A6})$$

Finally it is also straightforward to get that

$$\begin{aligned} \langle \tilde{g}[k] \tilde{g}^*[k'] \rangle &= \frac{1}{2\pi} \int_{-\pi}^{\pi} d\phi \tilde{S}_{\tilde{g}}(\phi) \\ &\times w^*\left(\phi - k\frac{2\pi}{N}\right) w\left(\phi - k'\frac{2\pi}{N}\right), \end{aligned} \quad (\text{A7})$$

$$\begin{aligned} \langle \tilde{g}[k] \tilde{g}[k'] \rangle &= \frac{1}{2\pi} \int_{-\pi}^{\pi} d\phi \tilde{S}_{\tilde{g}}(\phi) \\ &\times w\left(\phi - k\frac{2\pi}{N}\right) w\left(\phi - k'\frac{2\pi}{N}\right). \end{aligned} \quad (\text{A8})$$

Thus, if k and k' are spaced enough that $w(\phi - k\frac{2\pi}{N})$ has no overlap with $w(\phi - k'\frac{2\pi}{N})$, then $\tilde{g}[k]$ and $\tilde{g}[k']$ are independent random variables.

The amount of spacing needed for all the conditions above to hold depends on the selected spectral window




FIG. 7. The spectral window for the Blackman-Harris plotted for different values of k , spaced by $k_o = 8$.




FIG. 8. The absolute value of the correlation coefficient defined in Eq. (A9), as a function of the DFT coefficient spacing Δk . Values are calculated for the Blackman-Harris window in the case of white noise.

$w[n]$. In general there exists a common value k_o for which, with good accuracy, Eq. (A4) holds if $|k| \geq \frac{k_o}{2}$, and $\tilde{g}[k]$ are independent random variables if $|k - k'| \geq k_o$. With the Blackman-Harris window used in the numerical calculations of the present paper, $k_o = 8$ (see Fig. 7).

The correlation coefficients resulting from Eq. (A7), may be calculated explicitly, for nearby DFT coefficients, assuming the noise is white. In Fig. 8 we report the values for

$$|\rho(\Delta k)| = \frac{|\langle \text{Re}\{\tilde{g}[k]\} \text{Re}\{\tilde{g}[k + \Delta k]\} \rangle|}{\sqrt{\langle |\text{Re}\{\tilde{g}[k]\}|^2 \rangle \langle |\text{Re}\{\tilde{g}[k + \Delta k]\}|^2 \rangle}} \\ = \frac{|\langle \text{Im}\{\tilde{g}[k]\} \text{Im}\{\tilde{g}[k + \Delta k]\} \rangle|}{\sqrt{\langle |\text{Im}\{\tilde{g}[k]\}|^2 \rangle \langle |\text{Im}\{\tilde{g}[k + \Delta k]\}|^2 \rangle}}. \quad (\text{A9})$$

For the first few values of Δk . In addition one can calculate that

$$\frac{|\langle \text{Re}\{\tilde{g}[k]\} \text{Im}\{\tilde{g}[k + \Delta k]\} \rangle|}{\sqrt{\langle |\text{Im}\{\tilde{g}[k]\}|^2 \rangle \langle |\text{Re}\{\tilde{g}[k + \Delta k]\}|^2 \rangle}} \sim 0. \quad (\text{A10})$$

Thus if noise does not vary too much over a range of order of k_o , a reduced spacing among coefficients may also be considered. In the case of the Blackman-Harris window, Fig. 8 suggests that coefficients could be taken every $\Delta k \sim 4$ and still be treated as basically independent, and that even for $\Delta k \sim 2-3$ the effect of the correlation may still be negligible.

APPENDIX B: ADDITIONAL FORMULAS

To get the result in Eqs. (18), and (19), it is sufficient to substitute $P(\vec{g}|\vec{\theta}, \vec{S})$ taken from Eq. (9) into Eq. (17) and perform the integral over S_j . We get

$$\langle S_j | \vec{g} \rangle = \frac{N_s}{N_s - m - 1} \frac{\int |\vec{g}[j, \vec{\theta}]|^2 \prod_{k \in Q} (|\vec{g}[k, \vec{\theta}]|^{-2})^{-(N_s - m)} P(\vec{\theta}) d\vec{\theta}}{\int \prod_{k \in Q} (|\vec{g}[k, \vec{\theta}]|^{-2})^{-(N_s - m)} P(\vec{\theta}) d\vec{\theta}} \\ = \frac{N_s}{N_s - m - 1} \langle |\vec{g}[j]|^2 \rangle, \quad (\text{B1})$$

$$\langle S_j^2 | \vec{g} \rangle = \frac{N_s^2}{(N_s - m - 1)(N_s - m - 2)} \\ \times \frac{\int |\vec{g}[j, \vec{\theta}]|^4 \prod_{k \in Q} (|\vec{g}[k, \vec{\theta}]|^{-2})^{-(N_s - m)} P(\vec{\theta}) d\vec{\theta}}{\int \prod_{k \in Q} (|\vec{g}[k, \vec{\theta}]|^{-2})^{-(N_s - m)} P(\vec{\theta}) d\vec{\theta}} \\ = \frac{N_s^2}{(N_s - m - 1)(N_s - m - 2)} \langle |\vec{g}[j]|^4 \rangle. \quad (\text{B2})$$

[1] S. Anza *et al.*, *Classical Quantum Gravity* **22**, S125 (2005).
 [2] L. Ferraioli, M. Hueller, and S. Vitale, *Classical Quantum Gravity* **26**, 094013 (2009).
 [3] G. Congedo, R. Dolesi, M. Hueller, S. Vitale, and W.J. Weber, *Phys. Rev. D* **88**, 082003 (2013).
 [4] S. Vitale, *Gen. Relativ. Gravit.* **46**, 1730 (2014).
 [5] P. Welch, *IEEE Trans. Audio Electroacoust.* **15**, 70 (1967).

[6] T. B. Littenberg, M. Coughlin, B. Farr, and W. M. Farr, *Phys. Rev. D* **88**, 084044 (2013).
 [7] C. Röver, A. Stroer, E. Bloomer, N. Christensen, J. Clark, M. Hendry, C. Messenger, R. Meyer, M. Pitkin, J. Toher, R. Umstätter, A. Vecchio, J. Veitch, and G. Woan, *Classical Quantum Gravity* **24**, S521 (2007).
 [8] C. Röver, R. Meyer, and N. Christensen, *Classical Quantum Gravity* **28**, 015010 (2011).

- [9] N. J. Cornish and T. B. Littenberg, *Phys. Rev. D* **76**, 083006 (2007).
- [10] A. Gelman, J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, *Bayesian Data Analysis, Third Edition* (Chapman and Hall, London, 2014).
- [11] D. B. Percival and A. T. Walden, *Spectral Analysis for Physical Applications* (Cambridge University Press, Cambridge, England, 1993).
- [12] H. Jeffreys, *Proc. R. Soc. A* **186**, 453 (1946).
- [13] N. Brandt, W. Fichter, M. Kersten, S. Lucarelli, and F. Montemurro, *Classical Quantum Gravity* **22**, S493 (2005).
- [14] F. Antonucci *et al.*, *Classical Quantum Gravity* **28**, 094006 (2011).
- [15] M. Nofrarias, C. Röver, M. Hewitson, A. Monsky, G. Heinzl, K. Danzmann, L. Ferraioli, M. Hueller, and S. Vitale, *Phys. Rev. D* **82**, 122002 (2010).
- [16] G. Congedo, L. Ferraioli, M. Hueller, F. De Marchi, S. Vitale, M. Armano, M. Hewitson, and M. Nofrarias, *Phys. Rev. D* **85**, 122004 (2012).
- [17] L. Ferraioli *et al.*, *Exp. Astron.* **37**, 109 (2014).
- [18] M. R. Adams and N. J. Cornish, *Phys. Rev. D* **89**, 022001 (2014).